



US006484225B2

(12) **United States Patent**  
**Sheikh et al.**

(10) **Patent No.:** **US 6,484,225 B2**  
(45) **Date of Patent:** **Nov. 19, 2002**

(54) **METHOD AND SYSTEM FOR MANAGING COMMUNICATIONS AMONG COMPUTER DEVICES**

4,835,737 A 5/1989 Herrig et al.  
4,949,245 A 8/1990 Martin et al.

(List continued on next page.)

(75) Inventors: **Tahir Q. Sheikh**, Fremont, CA (US);  
**Walter A. Wallach**, Los Altos, CA (US)

**FOREIGN PATENT DOCUMENTS**

JP 04 333 118 A 11/1992 ..... G06F/1/18  
JP 07 093 064 A 4/1995 ..... G06F/1/26

(73) Assignee: **Micron Technology, Inc.**, Boise, ID (US)

**OTHER PUBLICATIONS**

(\* ) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

Lyons, Computer Reseller News, Issue 721, pp. 61–62, Feb. 3, 1997, “ACC Releases Low–Cost Solution for ISPs.”  
M2 Communications, M2 Presswire, 2 pages, Dec. 19, 1996, “Novell IntranetWare Supports Hot Pluggable PCI from NetFRAME.”  
Rigney, PC Magazine, 14(17): 375–379, Oct. 10, 1995, “The One for the Road (Mobile–aware capabilities in Windows 95).”  
Shanley, and Anderson, PCI System Architecture, Third Edition, p. 382, Copyright 1995.

(21) Appl. No.: **09/854,236**

(22) Filed: **May 11, 2001**

(65) **Prior Publication Data**

US 2001/0025329 A1 Sep. 27, 2001

**Related U.S. Application Data**

(62) Division of application No. 09/048,909, filed on Mar. 26, 1998, now Pat. No. 6,421,746.

(51) **Int. Cl.**<sup>7</sup> ..... **G06F 13/14**

(52) **U.S. Cl.** ..... **710/310; 710/309**

(58) **Field of Search** ..... 710/48, 107, 306,  
710/309, 310, 52, 308

(56) **References Cited**

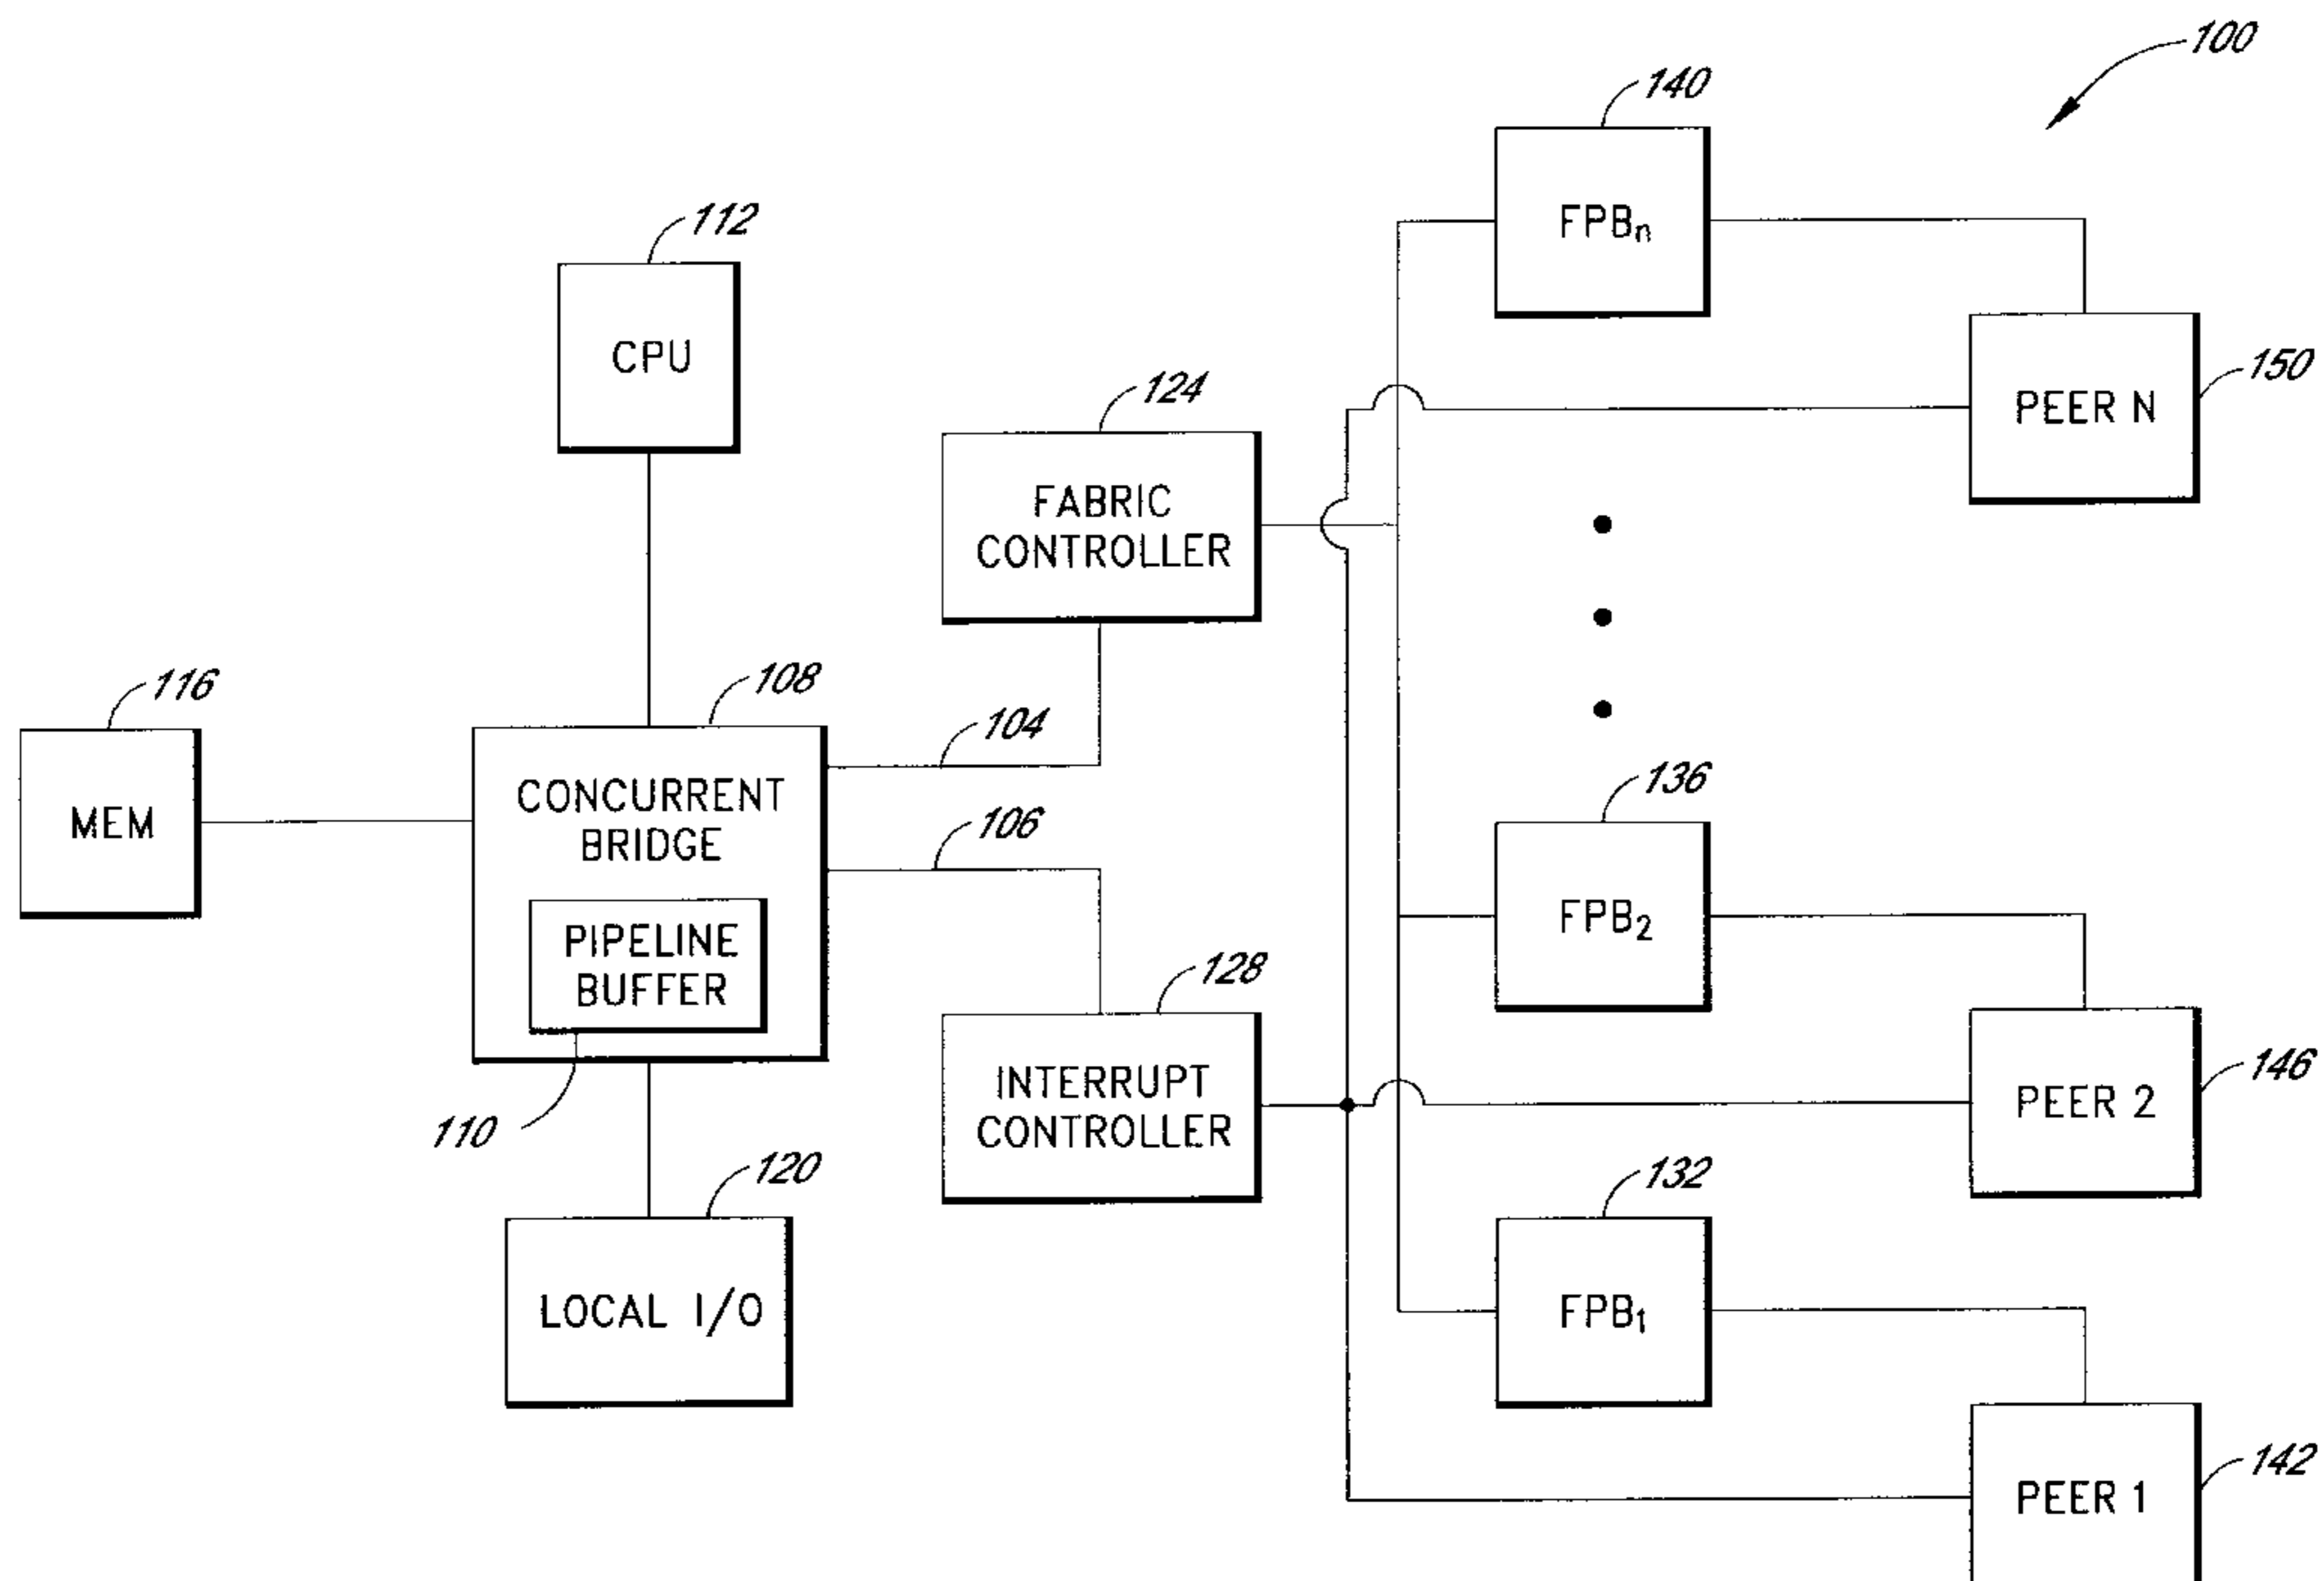
**U.S. PATENT DOCUMENTS**

4,449,182 A 5/1984 Rubinson et al.

(57) **ABSTRACT**

A method and system for managing communications among computer devices without involving central processor units of computer systems when it is determined that involving a central processor unit is unnecessary. The method employs a controller to manage communications among peer and host devices. With this method, congestion due to control and data traffic is minimized and a more efficient operation of central processor units is achieved.

**14 Claims, 3 Drawing Sheets**



U.S. PATENT DOCUMENTS

4,968,977 A	*	11/1990	Chinnaswamy et al. . 340/825.8	5,740,378 A	4/1998	Rehl et al.
4,999,787 A		3/1991	McNally et al.	5,747,889 A	5/1998	Raynham et al.
5,210,855 A		5/1993	Bartol	5,748,426 A	5/1998	Bedingfield et al.
5,265,098 A		11/1993	Mattson et al.	5,754,797 A	5/1998	Takahashi
5,269,011 A		12/1993	Yanai et al.	5,761,033 A	6/1998	Wilhelm
5,272,584 A		12/1993	Austruy et al.	5,761,045 A	6/1998	Olson et al.
5,317,693 A		5/1994	Cuenod et al.	5,761,454 A	*	6/1998 Adusumilli et al. .... 710/105
5,317,747 A		5/1994	Mochida et al.	5,764,924 A	6/1998	Hong
5,329,625 A		7/1994	Kannan et al.	5,764,968 A	6/1998	Ninomiya
5,337,413 A		8/1994	Lui et al.	5,765,198 A	6/1998	McCrocklin et al.
5,353,415 A	*	10/1994	Wolford et al. .... 711/115	5,768,541 A	6/1998	Pan-Ratzlaff
5,357,614 A	*	10/1994	Pattisam et al. .... 710/68	5,768,542 A	6/1998	Enstrom et al.
5,386,567 A		1/1995	Lien et al.	5,781,767 A	7/1998	Inoue et al.
5,426,740 A	*	6/1995	Bennett ..... 710/126	5,781,798 A	7/1998	Beatty et al.
5,446,910 A		8/1995	Kennedy et al.	5,784,576 A	7/1998	Guthrie et al.
5,471,634 A	11/1995		Giorgio et al.	5,790,831 A	8/1998	Lin et al.
5,483,419 A	1/1996		Kaczeus, Sr. et al.	5,793,987 A	8/1998	Quackenbush et al.
5,493,574 A	2/1996		McKinley	5,793,992 A	*	8/1998 Steele et al. .... 710/113
5,493,666 A	2/1996		Fitch	5,794,035 A	8/1998	Golub et al.
5,517,646 A	5/1996		Piccirillo et al.	5,796,185 A	8/1998	Takata et al.
5,530,810 A	6/1996		Bowman	5,796,981 A	8/1998	Abudayyeh et al.
5,539,913 A	7/1996		Furuta et al.	5,798,828 A	8/1998	Thomas et al.
5,555,510 A	9/1996		Verseput et al.	5,799,036 A	8/1998	Staples
5,564,024 A	10/1996		Pemberton	5,802,269 A	9/1998	Poisner et al.
5,568,610 A	10/1996		Brown	5,802,393 A	9/1998	Begun et al.
5,579,491 A	11/1996		Jeffries et al.	5,802,552 A	9/1998	Fandrich et al.
5,581,712 A	12/1996		Herrman	5,805,834 A	9/1998	McKinley et al.
5,586,250 A	12/1996		Carbonneau et al.	5,809,224 A	9/1998	Schultz et al.
5,588,121 A	12/1996		Reddin et al.	5,812,757 A	9/1998	Okamoto et al.
5,588,144 A	12/1996		Inoue et al.	5,812,858 A	9/1998	Nookala et al.
5,606,672 A	2/1997		Wade	5,815,117 A	9/1998	Kolanek
5,608,876 A	3/1997		Cohen et al.	5,822,547 A	10/1998	Boesch et al.
5,615,207 A	3/1997		Gephardt et al.	5,838,935 A	*	11/1998 Davis et al. .... 710/129
5,632,021 A	5/1997		Jennings et al.	5,870,567 A	*	2/1999 Hausauer et al. .... 710/112
5,638,289 A	6/1997		Yamada et al.	5,878,237 A	*	3/1999 Olarig ..... 710/128
5,644,470 A	7/1997		Benedict et al.	5,890,002 A	*	3/1999 Li et al. .... 710/260
5,644,731 A	7/1997		Lienres et al.	5,909,568 A	6/1999	Nason
5,651,006 A	7/1997		Fujino et al.	5,911,779 A	6/1999	Stallmo et al. .... 714/6
5,652,832 A	7/1997		Kane et al.	5,918,057 A	*	6/1999 Chou et al. .... 710/260
5,664,119 A	9/1997		Jeffries et al.	5,930,358 A	7/1999	Rao ..... 380/4
5,680,288 A	10/1997		Carey et al.	5,953,538 A	*	9/1999 Duncan et al. .... 710/22
5,696,970 A	12/1997		Sandage et al.	5,964,855 A	10/1999	Bass et al.
5,721,935 A	2/1998		DeSchepper et al.	5,983,349 A	11/1999	Kodama et al. .... 713/200
5,726,506 A	3/1998		Wood			

\* cited by examiner

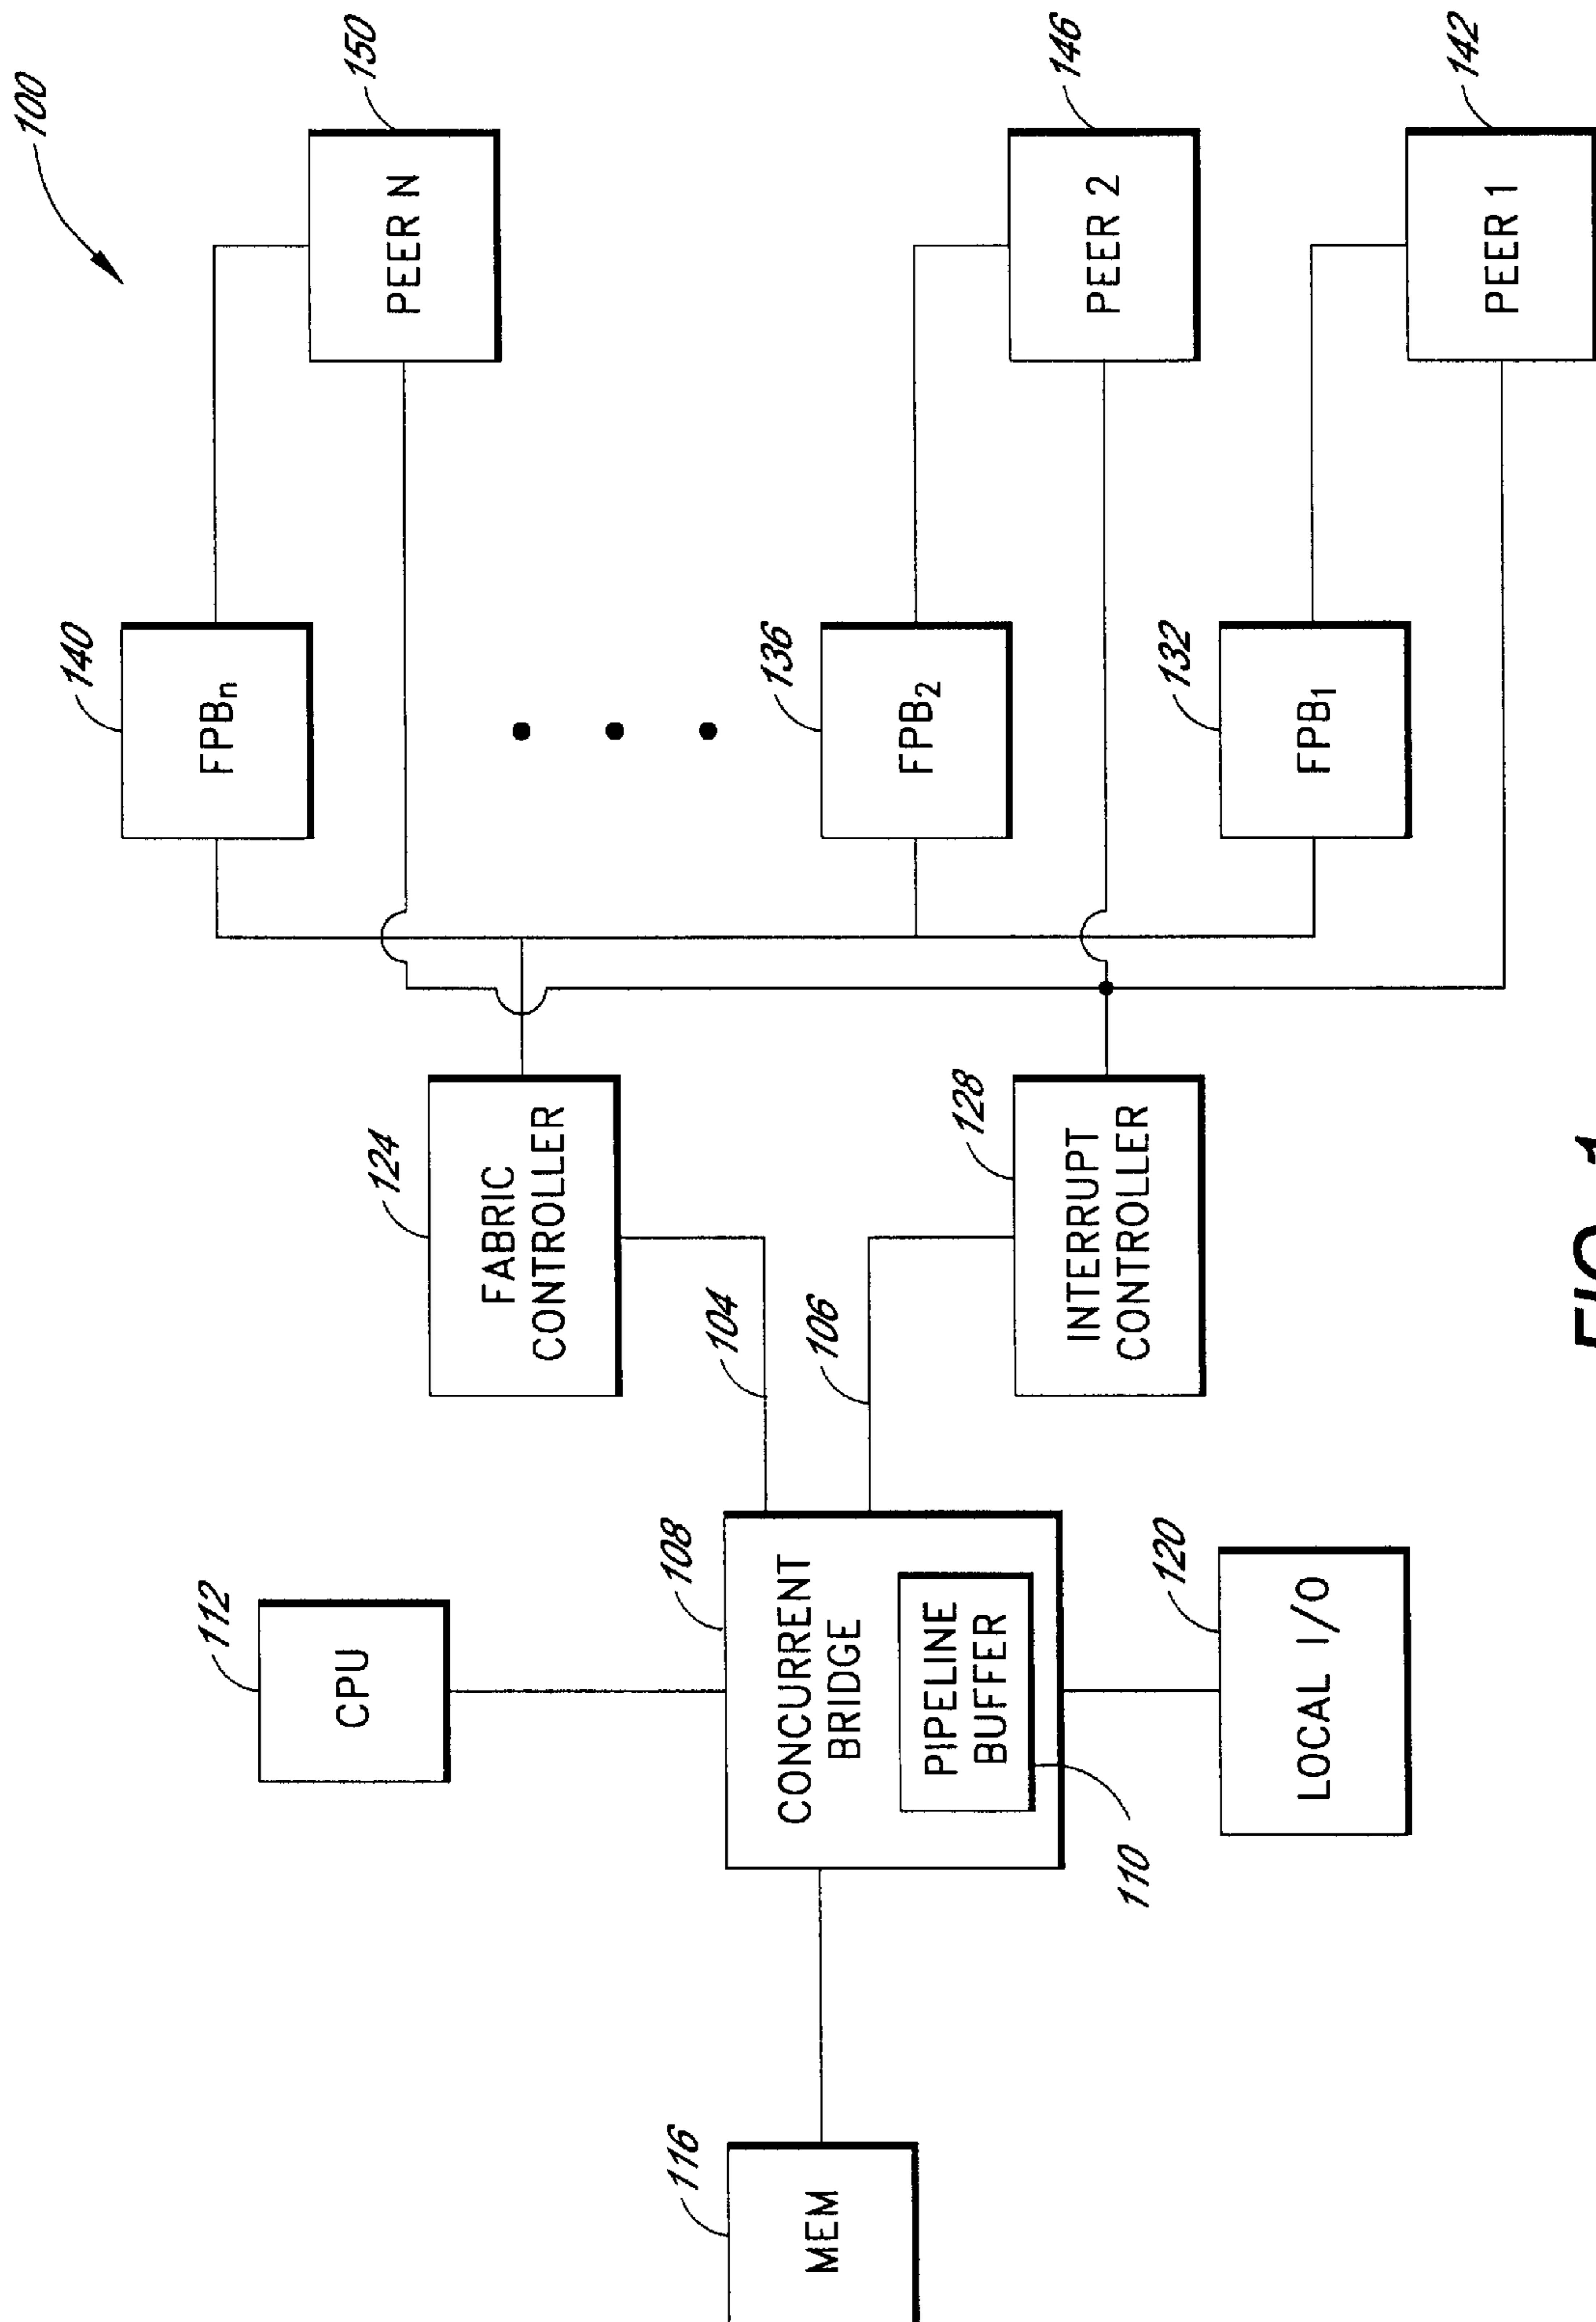


FIG. 1

FIG. 2

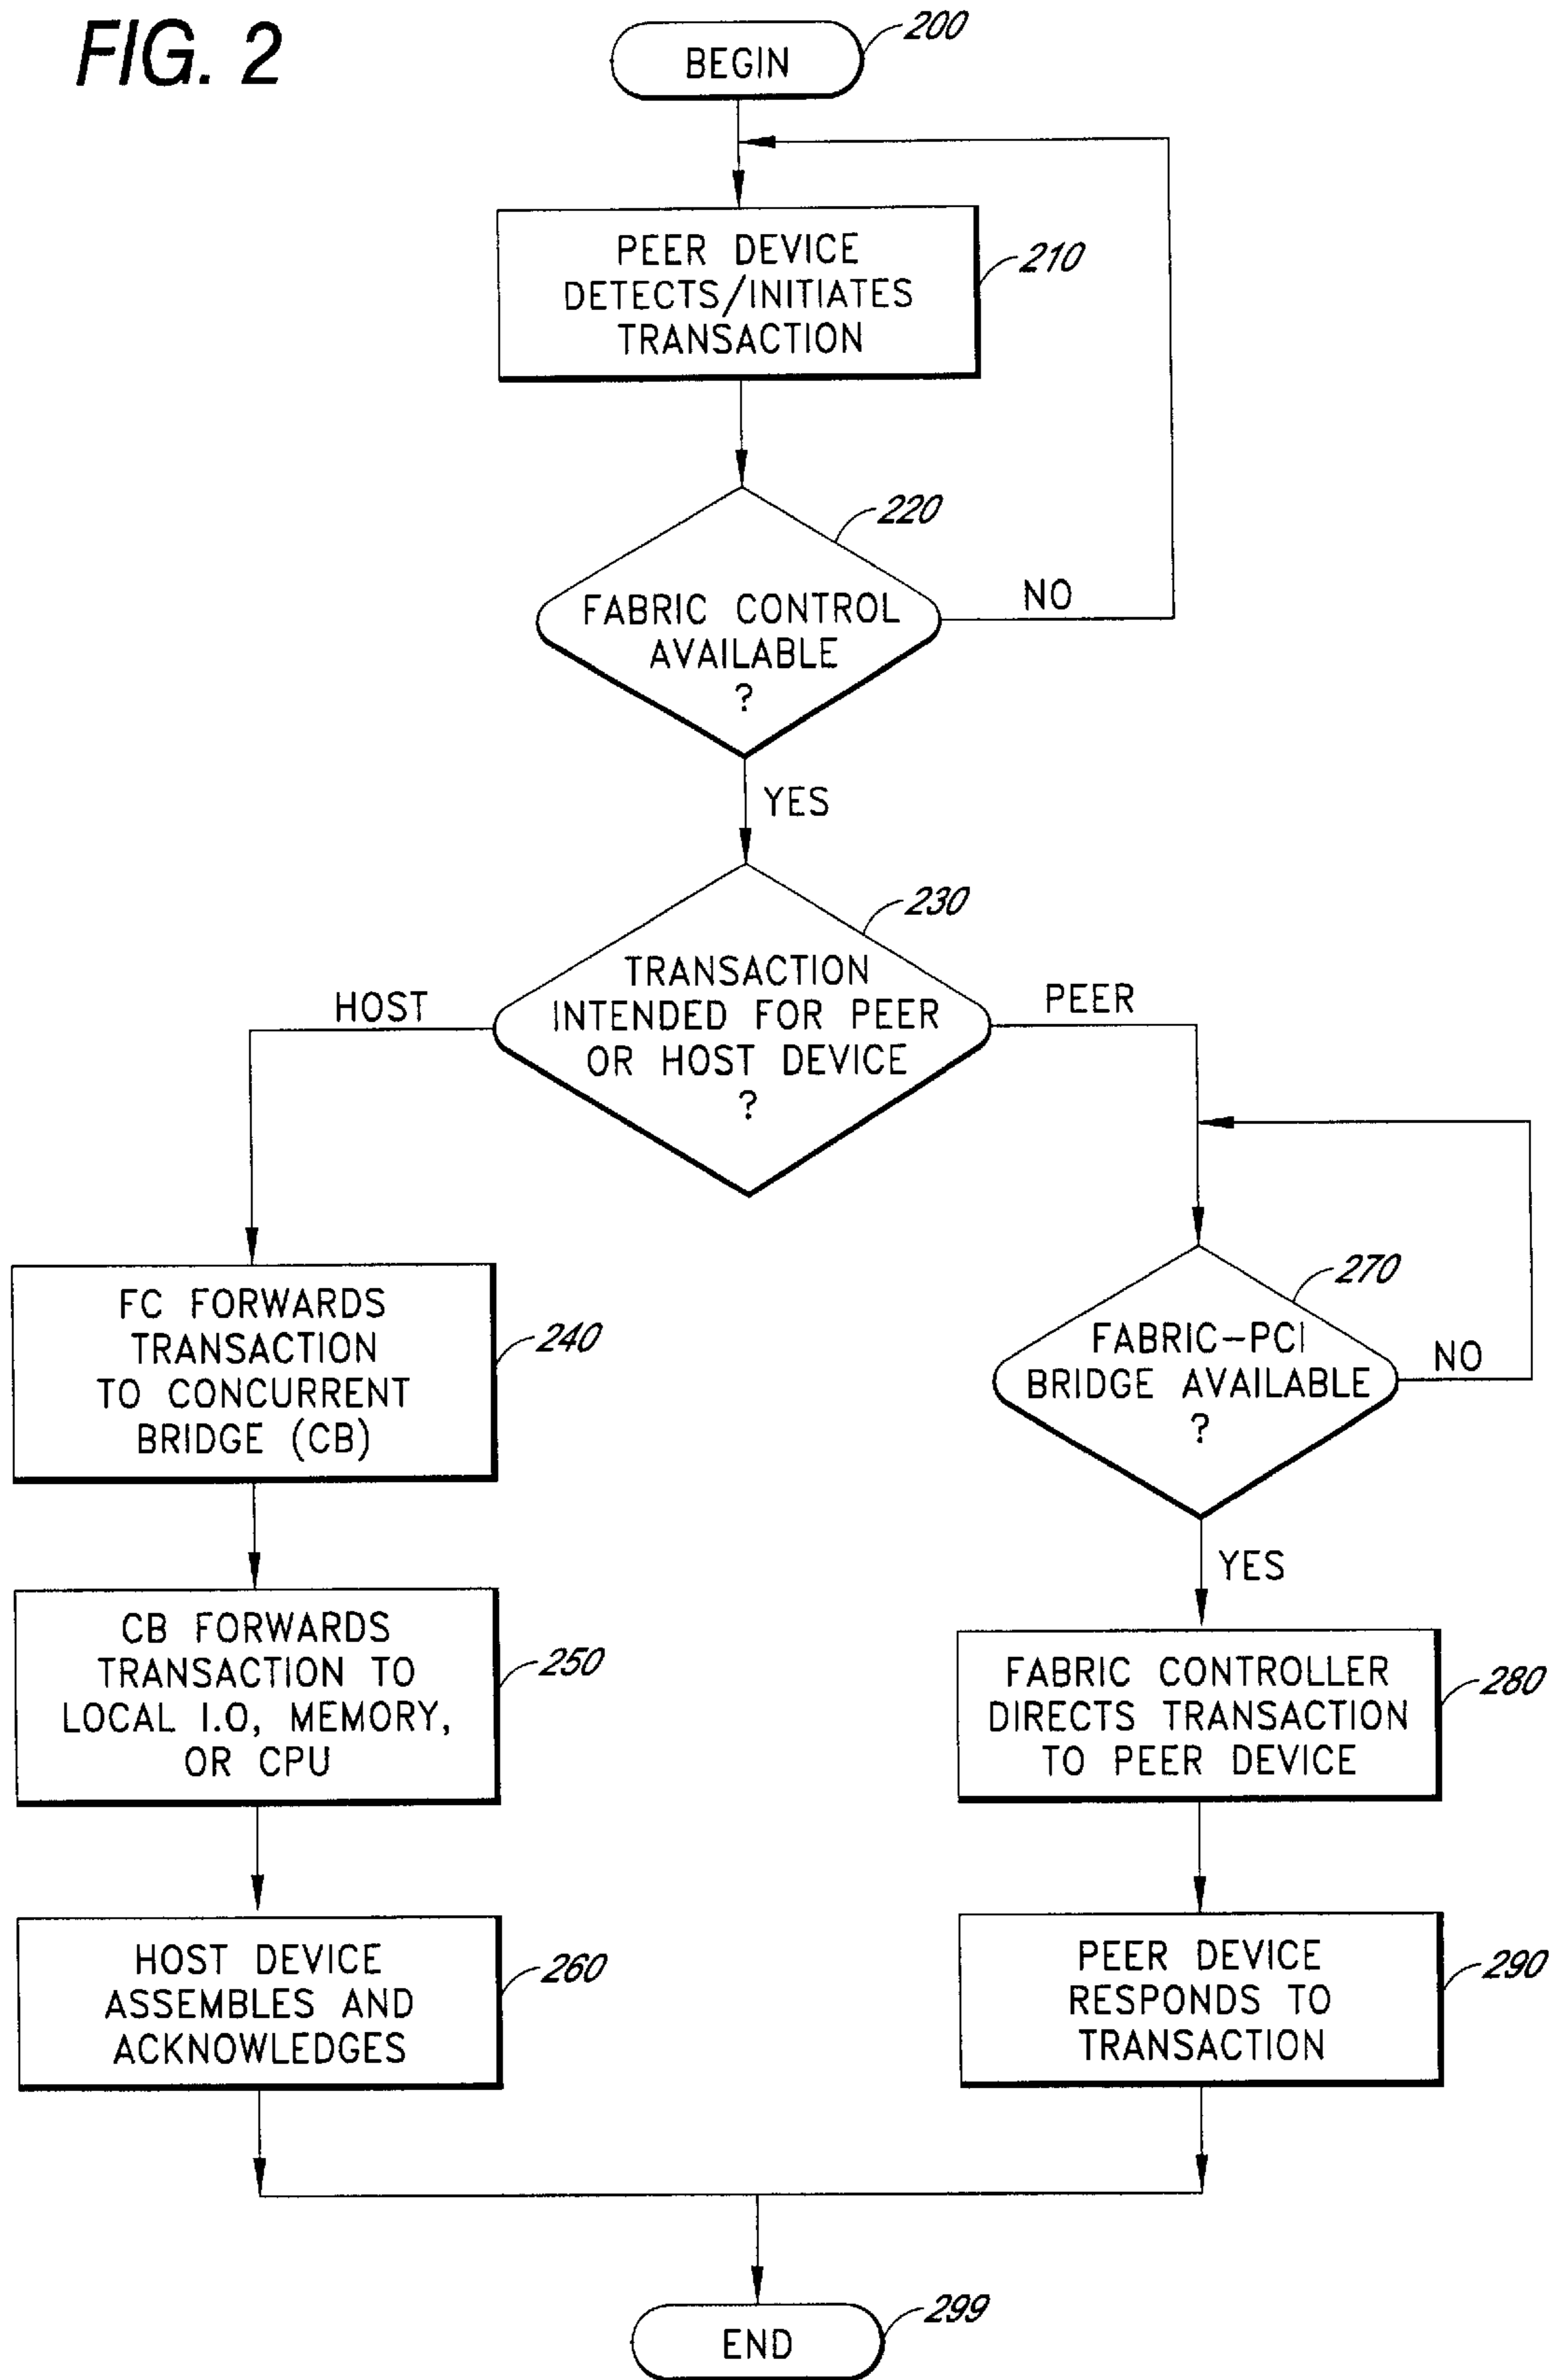
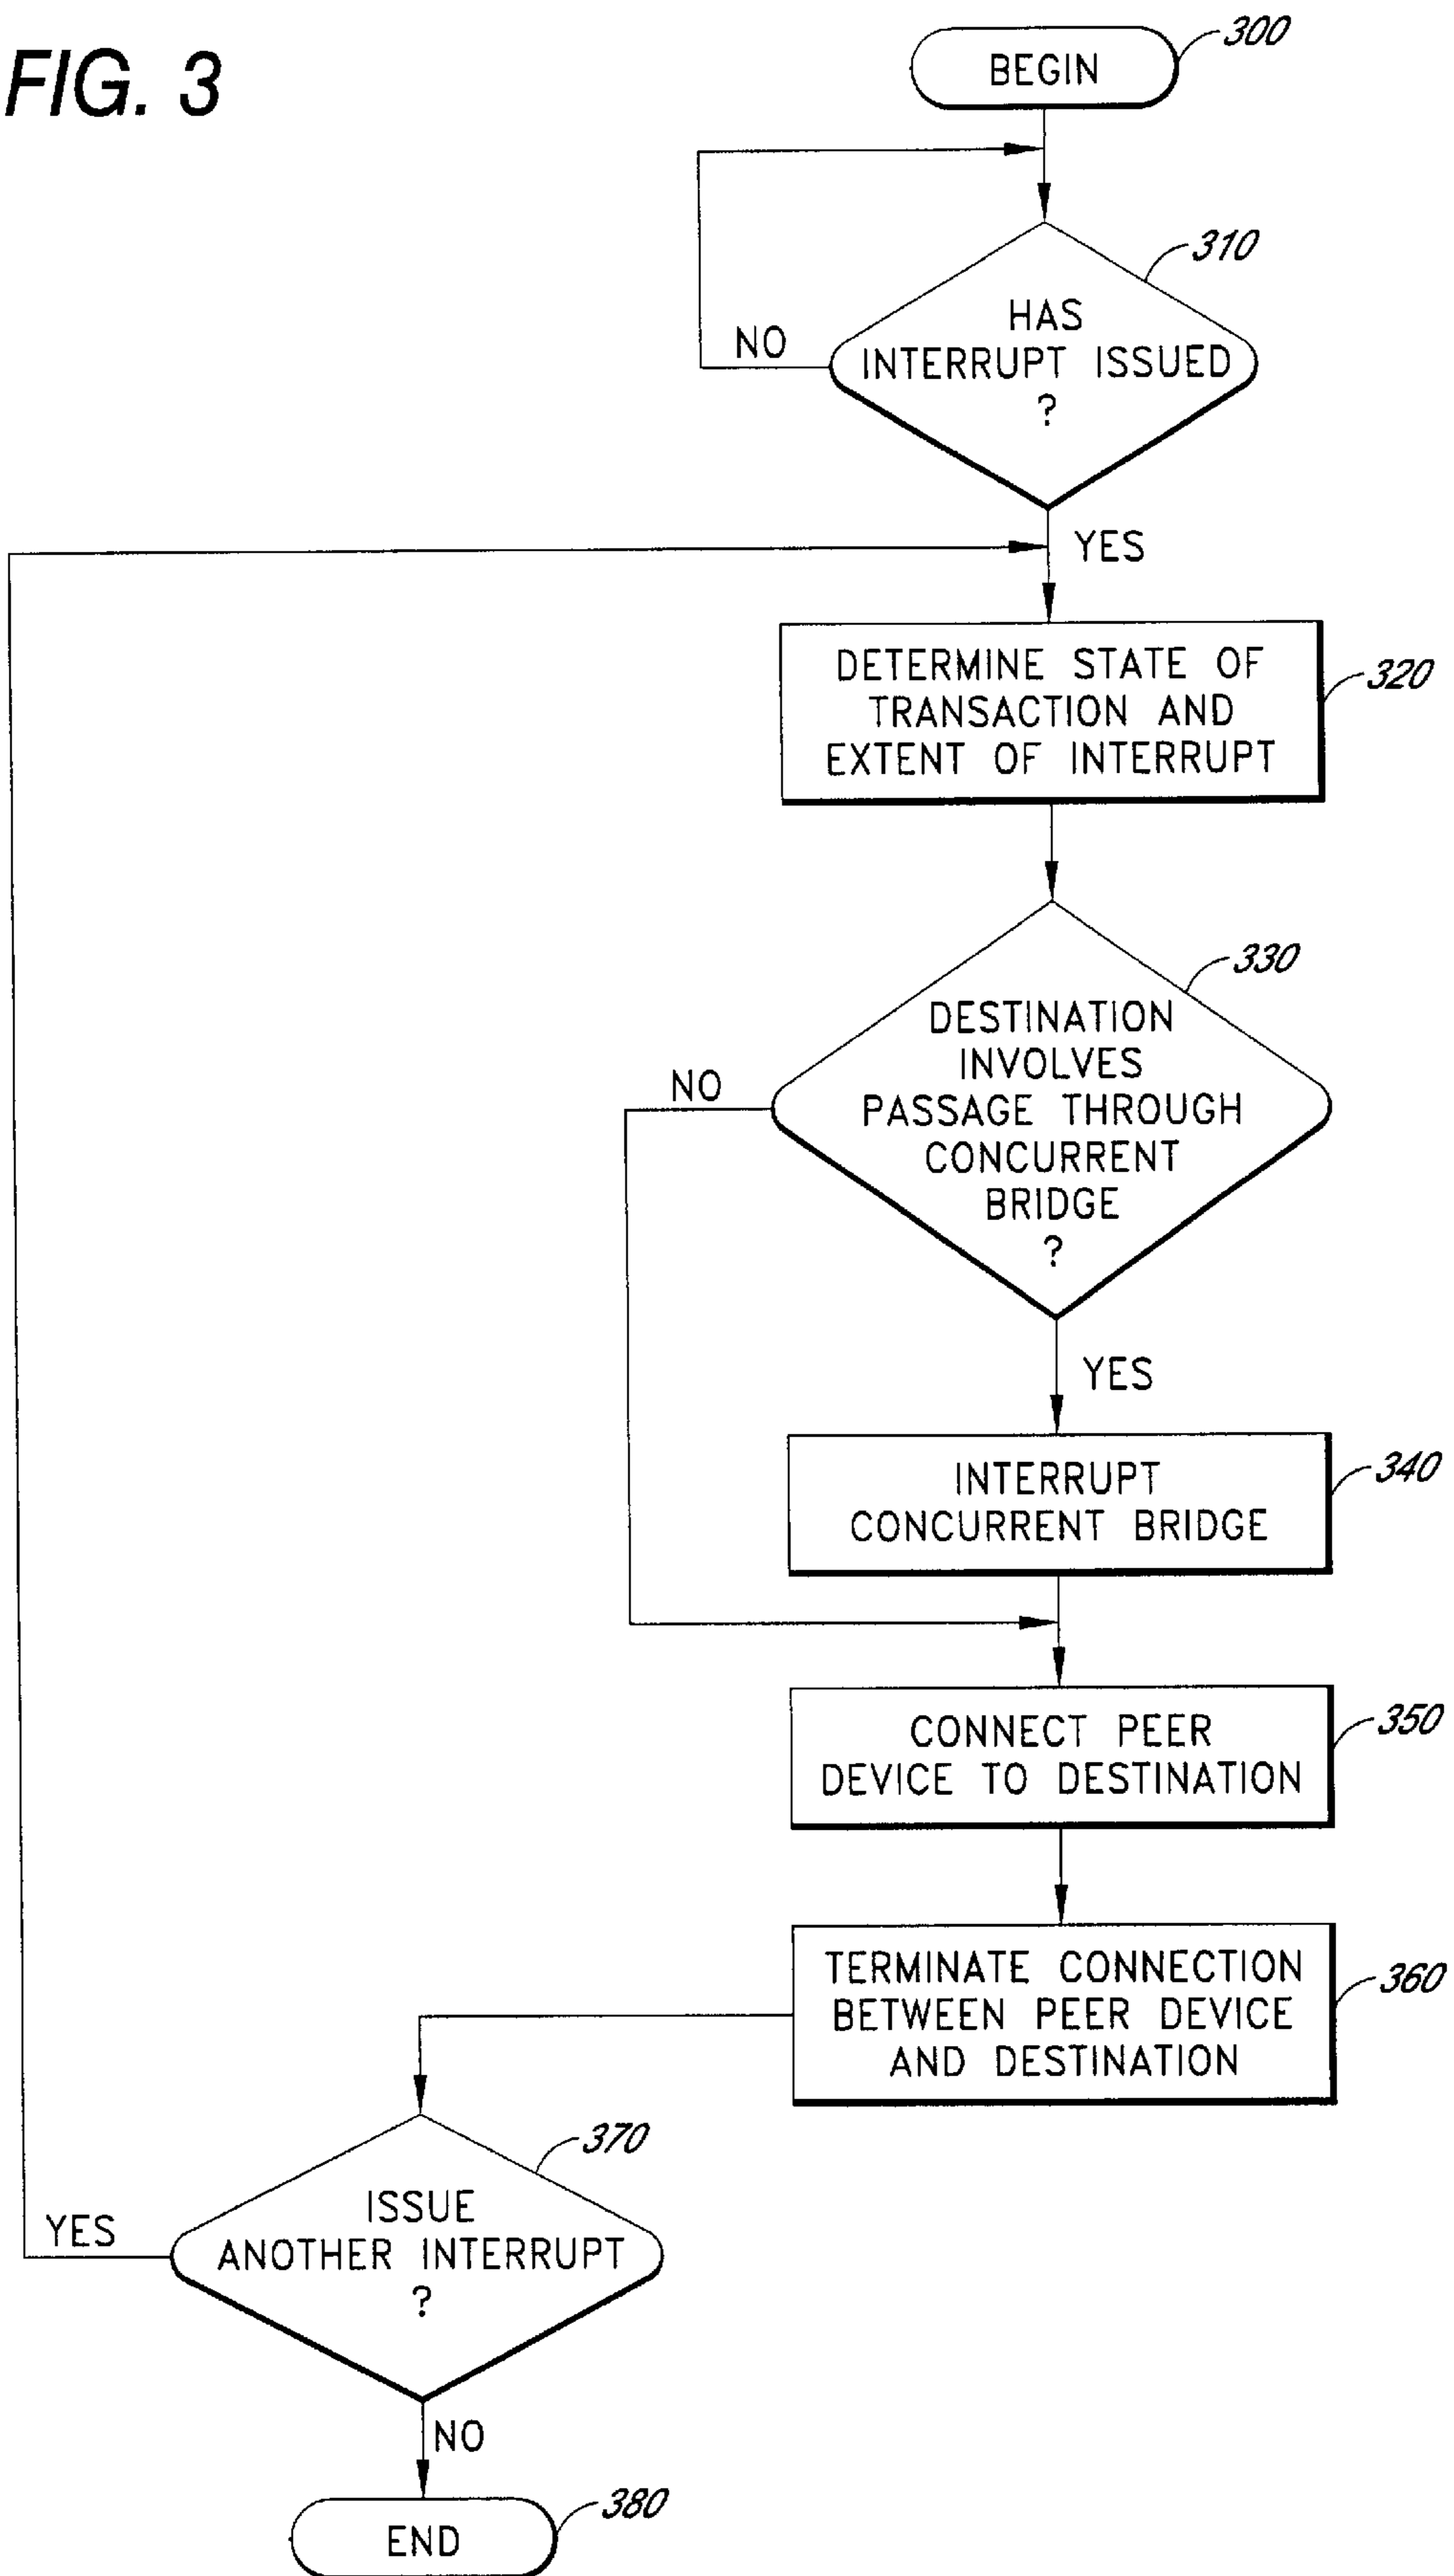


FIG. 3



## METHOD AND SYSTEM FOR MANAGING COMMUNICATIONS AMONG COMPUTER DEVICES

### RELATED APPLICATIONS

This application is a divisional of U.S. patent application entitled "METHOD OF DATA AND INTERRUPT POSTING FOR COMPUTER DEVICES" having Application Ser. No. 09/048,909, filed on Mar. 26, 1998 now U.S. Pat. No. 6,421,746.

The subject matter of U.S. Patent Application entitled SYSTEM FOR DATA AND INTERRUPT POSTING FOR COMPUTER DEVICES, filed on Mar. 26, 1998, Application Ser. No. 09/048,818, and having attorney Docket No. MNFRAME.068A is related to this application.

### BACKGROUND OF THE INVENTION

#### 1. Field of the Invention

The invention relates generally to data processing within information processing systems. More particularly, this invention relates to inter-device communication within a computer system.

#### 2. Description of the Related Art

Information processing systems, such as personal computers (PCs), have virtually become an inseparable part of many people's daily activities. These systems process an enormous amount of information in a relatively short time. To perform these sophisticated tasks, a computer system typically includes a central processor, memory modules, various system and bus control units, and a wide variety of peripheral data input/output (I/O) and storage devices. These computer components communicate using control and data signals having various data rates and signal protocols over multiple system buses.

Examples of such system buses include a peripheral component interconnect ("PCI") bus, a scaleable coherent interface ("SCI") bus, and a high performance parallel interface ("HIPPI") bus. The PCI bus is a 32-bit or 64-bit bus with multiplexed address and data lines. The bus is intended for use as an interconnect mechanism between highly integrated peripheral controller components, peripheral add-in boards, and processor/memory devices. In some applications, the SCI bus uses point-to-point links and a packet protocol to support 64-bit physical addresses. The upper 16 bits of the 64-bit address specify a node number and the lower 48 bits of the 64-bit address specify an offset address. The SCI bus uses coaxial cables over medium distances (e.g., 10's of meters) and fiber optics over long distances (e.g., 10 km) to provide unidirectional point-to-point signaling, from a transmitting device (i.e., transmitter) to a receiving device (i.e., receiver), to simulate a bus. The SCI bus supports read and write transactions among the various devices within a computer system. A transaction includes request and response subactions. The request subaction transfers an address and a command (read or write), whereas the response subaction returns status. For a write transaction, data are included within the request packet. For a read transaction, data are included within the response packet. For a compound transaction (e.g., fetch and add), data are included within the request and response packets.

The HIPPI protocol supports bus communication over a simplex channel (point-to-point link) for transferring data in one direction. In some applications, the HIPPI bus uses a parallel data path to provide communication at 800 Mbps with a 32-bit data bus, and 1.6 Gbps with a 64-bit data bus.

The HIPPI bus performs data transfers and flow control in increments of bursts, with each burst nominally containing 256 words (i.e., 1024 or 2048 bytes). The HIPPI bus provides error detection by using byte parity on the data bus, and immediately following each burst of data with a length/longitudinal redundancy checkword (LLRC). HIPPI framing protocol (FP) defines the framing for packets that will be sent over a HIPPI connection. Basically the HIPPI-FP standard splits a packet in three areas: Header\_Area, D1\_Area, and D2\_Area. Each of these areas starts and ends on a 64-bit boundary. The Header\_Area defines the sizes and offsets of the D1\_Area and D2\_Area. The D1\_Area contains control information and the D2\_Area contains data associated with the control information.

Despite the transfer power of these communication protocols, data and control traffic among computer devices is still prevalent. Bottlenecks of data and control traffic among central processing units ("CPUs"), memory devices, and external media all adversely affect processing speeds and efficiency rates of computer systems. Data and control transactions are often limited to a common path used by all devices in the system. For instance, data traffic for devices on various input/output ("I/O") buses travels through the host processor bus. Additionally, all communications among peer devices travel through the host processor bus. Peer devices on the PCI bus may include one or more of the following: an audio card, a motion video card, a small computer system interface (SCSI) card, a graphics card, or other PCI-PCI bridges. For each transaction, a peer device may issue one or more interrupts to the processor to communicate to another device in the system. The frequency of interrupts results in unnecessary and often excessive data traffic on the host processor bus. More importantly, the involvement of the CPU in the management of these transactions slows computer processing speeds significantly.

Several attempts have been made in the field to resolve the bottleneck of traffic resulting from the above-described common path. Some of these attempts include employing data paths having higher data rate capacity, or widening data path bandwidths to support higher data throughput on the bus. These solutions, however, are often costly and, more importantly, limited by the capacity of the employed data path. Therefore, there is a need in the computer technology to manage device interrupts more effectively. The solution should provide a more efficient utility of CPUs while continuing to meet the demands of increasing control and data traffic.

### SUMMARY OF THE INVENTION

To overcome the limitations of the related art, the invention provides a method of posting data and interrupt transactions for devices and local subsystems in a computer system. A local subsystem may include one or more peer devices. The invention provides a fabric controller, a concurrent bridge, and an interrupt controller to alleviate the need of burdening the CPU with every transaction in the system. Accordingly, unnecessary control and data flow through the host processor bus is minimized.

### BRIEF DESCRIPTION OF THE DRAWINGS

The above and other aspects, features and advantages of the invention will be better understood by referring to the following detailed description, which should be read in conjunction with the accompanying drawings, in which:

FIG. 1 is a functional block diagram of a computer system employing one embodiment of the invention.

FIG. 2 is a flow chart describing the decisional steps of one embodiment of the Fabric Controller.

FIG. 3 is a flow chart describing the decisional steps of one embodiment of the interrupt controller.

#### DETAILED DESCRIPTION OF THE INVENTION

The invention provides a method of posting data and interrupt transactions for devices in a computer system. The method employs a fabric controller, a concurrent bridge, and an interrupt controller to alleviate the need of burdening the CPU with every transaction in the system. Accordingly, unnecessary control and data flow through the host processor bus is minimized. Additionally, by directing necessary transactions to the CPU, concurrent data and control transactions in a single system are supported.

FIG. 1 shows a functional block diagram of a computer system employing one embodiment of the invention. As shown in FIG. 1, a computer system 100 comprises a plurality of host devices communicating via a concurrent bridge 108 using standard I/O data buses. These host devices include, for example, a central processing unit ("CPU") 112, one or more memory units 116, and a local input/output ("I/O") interface 120 for connecting one or more local I/O devices. The invention is implemented independently of the bus protocol used. Accordingly, the concurrent bridge (CB) bus 104 may be one of a variety of bus protocols which are well known in the art. For example, in one embodiment, the CB bus 104 may be a scalable coherent interface ("SCI") bus, or a high performance parallel interface ("HIPPI") bus. A fabric controller 124 is connected to the CB 108 via a data port (not shown) to manage the flow of transaction requests among peer devices, and among peer and host devices. Additionally, an interrupt controller 128 is connected to the CB 108 via a control port 106 to manage the flow of interrupt activity among peer devices, and among peer and host devices. The design of the CB 108, fabric controller 124, and interrupt controller 128 may be based on an application specific integrated circuit (ASIC).

The CB 108 includes four data ports to connect the CPU 112, memory 116, local I/O 120, and the fabric controller 124. The CB 108 further includes a control port 106 to connect the interrupt controller 128 to other host devices. The CB 108 establishes communications for up to two links simultaneously. As used in the patent document, a link refers to an internal connection between two ports within the CB 108. Hence, for example, the CPU 112 may communicate with the Local I/O 120, while the fabric controller 124 accesses the memory 116 simultaneously. In addition to its ability to establish concurrent links, the CB 108 includes an arbiter which coordinates access by competing devices to same resources. The CB 108 may utilize an internal pipeline buffer 110 to coordinate access to the same resource. Hence, for example, if the fabric controller 124 is communicating to the memory 116 and the CPU 112 requests access to write into the memory 116 at the same time, the CB 108 allows the CPU 112 to write into the pipeline buffer 110 of the CB 108. After the fabric controller 124 completes its communication with the memory 116, the CB writes data stored in its pipeline buffer 110 into memory 116. Hence, the CB 108 provides virtual access by competing devices to the same resource simultaneously.

In addition to the main host bus, computer systems typically include other buses to support communication among peripheral devices, and between the CPU 112 and peripheral devices. One very common bus is the peripheral

component interconnect ("PCI") bus which supports communication by PCI devices to host and other devices in the system. A plurality of fabric-PCI bridges ("FPBs") provide bus protocol conversion to connect PCI buses to the CB bus 104. In this embodiment, a FPB1 132, FPB2 136, and FPBn 140 are connected to the CB bus 104 to provide communication for a plurality of PCI devices. As noted above, typical PCI devices ("peer devices") include an audio card, a motion video card, a local area network (LAN) interface, a SCSI card, an expansion bus interface, a graphics card, or other PCI-PCI bridges. As shown in FIG. 1, peer devices resident on PCI buses include Peer1 142, Peer2 146, and PeerN 150. Data and control traffic transmitted by peer and host devices travel through, and under the management of, the fabric controller 124. Interrupt traffic transmitted by peer and host devices travels through, and under the management of, the Interrupt controller 128.

FIG. 2 is a functional flow chart describing the decisional steps of one embodiment of the fabric controller 124. The fabric controller 124 may be a processor-based unit which includes hardware and software in its design. The computer hardware architecture shown in FIG. 1 may be used as the basis for applying the decisional steps as executed by the fabric controller 124.

Typically, transaction requests by peer and host devices are issued continuously in the computer system 100. When a peer or host device is not issuing, receiving, or processing a transaction, the device is in an idle state as indicated at the beginning of the process at step 200. There are at least three identifiable categories of transactions in the system 100. The first category is known as a "local" transaction which includes transactions being issued by and processed within the peer device itself. The second category is known as a "global peer" transaction which includes transactions being issued by a peer device to one or more other peer devices for further action. A third category is known as a "global" transaction which includes transactions transferred between one or more peer devices and one or more host devices. More particularly, examples of a global transaction include a transfer between the CPU 112 and Peer 1 142, the memory 116 and Peer 1 142, and the local I/O 120 and Peer 1 142.

To perform its sophisticated management functions, the fabric controller 124 monitors the issuance, transfer, and completion of transactions using the following process. At step 210, a peer device detects or issues a transaction. The form of a transaction depends on the bus protocol employed among peer devices. In some bus protocols, the transaction command is communicated in the form of a packet. The packet includes, among other things, a source address, a destination address, a transaction address, a transaction type, one or more status bits, and one or more error correction bits (e.g., cyclic redundancy checksum CRC). A peer device (e.g., Peer1 142 of FIG. 1) may detect a transaction command which is received from another device or, alternatively, issued by Peer1 142 itself. At step 220, Peer1 142 checks for the availability of the fabric controller 124 for managing the transaction command being issued or transferred. Typically, Peer1 142 sends a synchronizing packet to establish a handshake with the fabric controller 124, and waits for an acknowledgment packet from the fabric controller 124. If the fabric controller 124 is not available, then Peer1 142 waits for the fabric controller 124 to send the acknowledgment packet to Peer1 142. The waiting arises when the fabric controller 124 is managing other transaction commands from other devices in the system. When the fabric controller 124 becomes available, then at step 230, the fabric controller 124 issues an acknowledg-

ment packet to and receives the transaction command from Peer1 142 via the FPB1 132. As noted above, the transaction command may be a read, write, or a compound subaction. The fabric controller 124 determines the intended destination of the transaction command pursuant to the destination address field in the packet.

If the transaction command is intended for a host device, then at step 240, the fabric controller 124 forwards the transaction command to the CB 108 (FIG. 1) for further action. At step 250, the CB 108, in turn, forwards the transaction command to its intended destination (e.g., CPU 112, memory unit 116, or local I/O interface 120) for processing. At step 260, the recipient host device returns a response packet to the issuing device to acknowledge that the transaction command has been received for processing. If, on the other hand, the fabric controller 124 determines in step 230 that the transaction command is intended for another peer device, then the fabric controller 124 moves to step 270. At step 270, the fabric controller 124 checks for the availability of the fabric-PCI bridge (e.g., FPB2 136) to which the intended peer device (e.g., Peer2 146) is connected. If the FPB2 136 is not available, the fabric controller 124 waits until it receives an acknowledgment packet from the FPB2 136. When the FPB2 136 becomes available, the FPB2 sends an acknowledgment packet to the fabric controller 124 and, at step 280, the fabric controller 124 directs the transaction command to Peer2 146 via the FPB2 136 for further action. At step 290, the recipient peer device responds to the transaction command by returning a response packet acknowledging receipt of the transaction request. By forwarding the transaction request directly to the intended peer device without involving the CPU 112, the possibility of bottle neck traffic on the CB bus 104 is minimized. Moreover, concurrent transactions among host devices and among peer devices are supportable. The process terminates at step 299.

FIG. 3 is a flow chart describing the decisional steps of one embodiment of the interrupt controller 128 (FIG. 1). As shown in FIG. 3, at step 300, a typical interrupt process commences by setting the interrupt controller 128 in a "watchdog" state and waiting for the issuance of interrupts by one or more peer devices. At step 310, the interrupt controller 128 determines if an interrupt has been issued by a peer device. If no interrupt has been issued, the interrupt controller 128 returns to its watchdog state as described in step 300. If an interrupt is detected then, at step 320, the interrupt controller 128 analyzes the state of the current transaction, which is being performed by the interrupt-issuing peer device ("source peer device"). Additionally, in response to the interrupt request by the source peer device, the interrupt controller 128 determines whether to interrupt a destination peer device (i.e., the peer device targeted by the source peer device) and/or interrupt the CPU 112 ("speculative interrupt").

In analyzing the state of the current transaction, the interrupt controller 128 determines whether the transaction is in its early stages of execution or nearing completion. The interrupt controller 128 may evaluate one or more factors to assess the state of the current transaction. The factors may include, among other things, the destination address, transaction address, one or more status bits, and type of transaction. Hence, for example, a source peer device (e.g., Peer1 142) may request to read certain data (e.g., an image) from the memory 116. If the transaction address specifies a block of data at the end of the image data for this type of transaction, the interrupt controller 128 determines that the read transaction is nearing completion. Alternatively, a sys-

tem operator may set a counter to a threshold rate to determine at which point (e.g., percentage) a transaction is nearing completion. For instance, if the image size is 8K, and Peer1 142 is reading the image data in blocks of 1K during each memory access, then the read transaction is nearing completion on the 8th access to memory. Therefore, if the interrupt controller determines that Peer1 142 is accessing the memory 116 for the 8th time, then the read transaction is nearing completion. If the transaction is nearing completion, the interrupt controller 128 may interrupt the CPU 112. Such interrupt may be necessary to prepare the CPU 112 for further action after the Peer1 142 completes executing its current transaction.

Moreover, in some instances, the interrupt controller 128 may interrupt the CPU 142 even when the transaction is not nearing completion. Such interrupt may be in response to an express request by the source peer device. Additionally, the interrupt controller 128 may interrupt the CPU 142 upon detecting an error in the transaction (e.g., a data overflow). In such case, the interrupt controller 128 interrupts the CPU 142 to take appropriate measures, e.g., instruct the source peer device to cancel or re-initiate the transaction.

At step 330, the interrupt controller 128 determines whether transaction packets sent by the source peer device to a destination device involve passage through the CB 108 (FIG. 1). Typically, a source peer device issues an interrupt command to communicate with another device in the system. More particularly, Peer1 142 issues an interrupt command indicating the address of the memory 116. Based on the address, the interrupt controller 128 determines whether interrupting the CB 108 is necessary to establish a data path between the Peer1 142 and the memory 116. Hence, if establishing a data path with the destination device involves passage through the CB 108, then at step 340, the interrupt controller 128 interrupts the CB 108 for this purpose.

On the other hand, if establishing a data path with the destination device does not involve passage through the CB 108, the interrupt controller 128 does not interrupt the CB 108. The process continues directly from step 330 to step 350. At step 350, the interrupt controller 128 establishes a data path between the Peer1 142 and the memory 116. The Peer1 142, in turn, reads the desired data from the memory 116. At step 360, the interrupt controller 142 terminates the data path between the peer device 142 and the memory 116. At step 370, the source device determines whether to issue another interrupt to communicate with another device. For instance, after reading and processing (e.g., expanding the image) the desired data, the peer device 142 may issue an interrupt to send out a request to write the processed (i.e., expanded image) data into another peer device (e.g., Peer2 146). Hence, if the Peer1 142 issues another interrupt to the interrupt controller 128, the process repeats at step 320. If, on the other hand, the Peer1 142 does not issue an interrupt to the interrupt controller 128, the process terminates at step 380.

In view of the foregoing, it will be appreciated that the invention overcomes the long-standing need for a method of managing data and interrupt commands issued by peer devices without the disadvantage of involving the central processor in every transaction. The invention ensures an effective utilization of central processors by minimizing unnecessary interruptions by other devices in a computer system. The invention may be embodied in other specific forms without departing from its spirit or essential characteristics. The described embodiment is to be considered in all respects only as illustrative and not restrictive. The scope of the invention is, therefore, indicated by the appended



claims rather than by the foregoing description. All changes which fall within the meaning and range of equivalency of the claims are to be embraced within their scope.

What is claimed is:

1. An apparatus for managing communications in a computer system having a central processor connected to a host bus, the apparatus comprising:

a plurality of communication ports configured to provide concurrent communications among a plurality of devices, wherein concurrent communications comprises a first communication between a first pair of the plurality of devices occurring simultaneously with a second communication between a second pair of the plurality of devices;

an arbiter operationally connected to at least one of the plurality of ports, the arbiter being configured to coordinate access by the plurality of devices to a particular device; and

a pipeline buffer operationally connected to at least one of the plurality of ports, the pipeline buffer being configured to provide access by one of the plurality of devices to a device that is currently unavailable.

2. The apparatus of claim 1, wherein the first or second communication is established via a connection between two of the plurality of ports.

3. The apparatus of claim 1, wherein the central processor is electrically connected to at least one of the communication ports, the central processor configured to control a communication initiated by at least one of the plurality of devices in the event that the communication necessitates involving the central processor.

4. The apparatus of claim 3, further comprising a memory and at least one input/output device, wherein the memory and the input/output device are operationally connected to the plurality of ports.

5. The apparatus of claim 1, further comprising an interrupt controller electrically connected to at least one of the plurality of devices and to at least one of the plurality of communication ports, the interrupt controller being configured to manage a communication initiated by at least one of the plurality of devices in the event that involvement by the central processor in the communication is unnecessary.

6. The apparatus of claim 1, wherein the pipeline buffer stores at least part of the communication initiated by a first one of the plurality of devices to a second one of the plurality of devices in the event that the second device is unavailable, and further wherein the pipeline buffer forwards the data to the second device when the second device becomes available.

7. A method of managing communications in a computer system having a central processor connected to a host bus, and having a controller connected to a nonhost bus, the method comprising:

initiating communication between a first device and a second device;

determining by the controller whether the communication between the first device and the second device necessitates involving the central processor, wherein the act

of determining by the controller comprises determining whether the first device has requested involvement in the communication by the central processor; and

communicating data between the first device and the second device under the control of the controller without involving the central processor if it is determined that involving the central processor is unnecessary.

8. The method of claim 7, wherein determining whether the communication necessitates involving the central processor comprises determining whether the communication is substantially complete.

9. The method of claim 7, wherein determining whether the communication necessitates involving the central processor comprises determining whether an error in the communication has occurred.

10. The method of claim 7, further comprising determining an intended destination of the communication.

11. The method of claim 10, wherein determining the intended destination comprises reading a destination address field in the communication.

12. The method of claim 7, wherein communicating data comprises communicating data between peer devices solely via the nonhost bus.

13. A system for managing communications in a computer system having a central processor connected to a host bus, and having a controller connected to a nonhost bus, the system comprising:

means for initiating communication between a first device and a second device;

means for determining whether the communication between the first device and the second device necessitates involving the central processor, wherein the determining means determines whether the first device has requested involvement in the communication by the central processor;

means for communicating data between the first device and the second device without involving the central processor if it is determined that involving the central processor is unnecessary.

14. In a computer system having a central processor connected to a host bus, and having a controller connected to a nonhost bus, a program storage device storing instructions that when executed by the computer perform a method of managing communications, the method comprising:

initiating communication between a first device and a second device;

determining by the controller whether the communication between the first device and the second device necessitates involving the central processor, wherein the act of determining by the controller comprises determining whether the first device has requested involvement in the communication by the central processor;

communicating data between the first device and the second device under the control of the controller without involving the central processor if it is determined that involving the central processor is unnecessary.